Report on Director's Review of Run IIb Upgrades

Executive Summary

Both experiments have made outstanding progress in moving toward the baseline review. Designs are clearly mature and all major aspects of the upgrades are supported by in-depth studies. The experimenters are to be congratulated for their efforts. We expect that this work will lead to an efficient and smooth upgrade process. We believe both experiments are very close to being ready for their baseline review.

It is recommended that both experiments use the same baseline luminosity conditions for Run IIb. Based on recent developments these appear to be 2×10^{32} /cm^2/sec with 396 ns bunch spacing for normal operation, and 4×10^{32} /cm^2/sec with 396 ns bunch spacing to demonstrate adequate headroom. These figures have been used in some cases but not all.

The two experiments presented quite different manpower estimates for the silicon detector projects. This difference has to be resolved before the baseline review. In addition, it would add a great deal of confidence to the silicon detector cost estimates if quantitative comparisons could be made with the costs of earlier silicon detectors built for these two experiments.

It will be very important for both experiments to build fully operational staves and layer 0 as soon as possible to reduce the technical and schedule risk in the silicon projects. The recent success of the SVX4 prototype will allow finalization of the hybrids which are on the critical path.

In many cases the performance of trigger systems at higher luminosities is estimated by linear extrapolation from lower rate conditions. More realistic simulations involving multiple events per bunch crossing should be done.

The plan for upgrading commercial processors during Run IIa, in preparation for Run IIb is rather different for the two experiments and is coupled to operating expenses for the two experiments. The difference in the upgrade plans should be justified or the plans reconsidered.

Some aspects of the TDRs would benefit from further attention. For CDF it would help to strengthen or expand the technical descriptions for the Central Preradiator, the Level-2 decision crate, and installation planning. For D0 there is no discussion of installation planning. There is also no overall table of contents.

1. Technical Subcommittee Report

Both experiments have made outstanding progress in moving toward the baseline review. Designs are clearly mature and all major aspects of the upgrades are supported by in-depth studies. The experimenters are to be congratulated for their efforts. We expect that this work will lead to an efficient and smooth upgrade process. We believe both experiments are very close to being ready for their baseline review.

1.1. The CDF and D0 Silicon Projects

1.1.1. Introduction

1.1.1.1. Both the CDF and D0 groups have submitted updated TDRs for the complete replacement of their silicon detectors. The maturity of the designs has increased greatly, and the scope of the proposed upgrades has not changed since the last review.

1.1.1.2. The CDF proposal is described in detail in 64 pages of TDR and supplemented with 600 lines of detailed schedule and 145 pages of WBS dictionary. The total project cost is about $18.2 M including escalation, contingency and overhead.

1.1.1.3. The D0 proposal is described in detail in 170 pages of TDR and supplemented with 1200 lines of detailed schedule and 280 pages of WBS dictionary. The total project cost is about $ 22.9 M including escalation, contingency and overhead.

1.1.1.4. The proponents should be congratulated on their successful efforts in preparing the reports and presentations. The committee also commends the cooperation between the two experiments on various technical issues.

1.1.2. Silicon sensors

1.1.2.1. Findings:

1.1.2.1.1. Both collaborations have selected high quality silicon sensors mainly from the same source.

1.1.2.1.2. D0 is considering the option of acquiring the sensors for layers 0 and 1 from another source and will make a decision based on performance after irradiations.

1.1.2.2. Comments:

1.1.2.2.1. The sensor's simple design and the fact that they are single-sided have lowered the risk greatly.

1.1.2.2.2. Layer 0 and 1 sensors performance dominate the impact parameter resolution and their quality is therefore crucial for overall performance.

1.1.2.3. Recommendations:

1.1.2.3.1. When deciding on the vendor for the silicon detectors of the inner layers, D0 should consider the increased risk of buying from a low-volume vendor with an unproven track record. Other factors to be considered include production yield, strip yield, and production stability, in addition to radiation performance.

1.1.3. Electronics and cables

1.1.3.1. Findings:

1.1.3.1.1. The SVX4 prototype chip has been received and tested, and it works very well. The experiments plan to fix some minor problems before production.

1.1.3.1.2. D0 has shown noise and pedestal results of a layer 1 module equipped with SVX4 chips and connected to ELMA detectors. This represents a major success on many fronts. CDF has also positively characterized the SVX4 chip mounted on a hybrid prototype, in addition to a full stave equipped with SVX3d chips. A full stave prototype with final components is being prepared.

1.1.3.1.3. Both D0 and CDF have made good progress on the analog cables for layer 0. In particular, CDF continues to develop a 50 μm pitch analog cable, while D0 now takes a conservative approach using 100 μm pitch stacked cable. The D0 approach greatly reduces the technical difficulties. D0 has already received two batches of prototypes with good yield. In addition, they have assembled a full layer 0 module and have studied the noise pick up problem. They have reproduced the present CDF layer 00 noise problem and found an effective grounding scheme to solve it.

1.1.3.1.4. Both experiments have experienced unexpected failures in some parts of their current detectors and have now tried to mitigate the effects of failure modes in their designs. This is particularly true in the wake of the CDF radiation accident of last March. This accident has been studied in detail, but the underlying mechanism is still not understood. As a partial countermeasure CDF has introduced a Priority Bypass Chip that would limit the effect of such a failure to a single hybrid instead a full readout chain.

1.1.3.2. Comments:

1.1.3.2.1. The success of the SVX4 chip is the single most important element of the presentations. It eliminates a major source of concern and puts the groups in the unusual condition of starting the project with an almost final readout IC in hand. The complete characterization of the SVX4 ASIC is very important before the next prototype is submitted. The different noise figures presented should be reconciled. The plan of acquiring from the November 2002 submission a large fraction of the chips ultimately needed, seems to be a good hedge against schedule risk.

1.1.3.2.2. The hybrid development is also on the critical path and great attention has to be paid to it.

1.1.3.2.3. The construction of module and stave final prototypes, with their potential for uncovering problems early in the construction phase, is a crucial step that the groups recognize and are pursuing at full speed.

1.1.3.2.4. The analog cables connecting the layer 0 strips to the hybrids are recognized as a technical risk by both the committee and the groups. The cables involve two major concerns: production yield and noise pick-up. Many experiments have experienced problems with low production yield of fine pitch flex cables, and CDF is experiencing noise pick-up problems with the analog flex cable in their current layer 00 detectors. The robustness of D0's stacked flex cable still needs to be evaluated, whereas the vendor providing CDF fine pitch cables has been known to have inconsistent quality in production, especially in the ability to wire-bond.

1.1.3.2.5. Although alternative design solutions for layer 0 might be possible, the current design is the one that has been proven for the CDF layer 00. Both groups are applying a great deal of R&D effort to ensure that the cables can be produced and that the layer 0 modules have adequate performance.

1.1.3.2.6. The unknown mechanism leading to failures in the CDF radiation accident raises concern about the sensitivity of the new design to high rates of radiation and to the possible existence of single points of failure.

1.1.3.3. Recommendations:

1.1.3.3.1. The committee encourages the groups to converge to a common technological solution for the layer 0 analog cables and to pursue this solution with multiple vendors. Noise suppression studies should continue in the short term. The groups should define a clear decision path and branch points to arrive at production. The quality of the cable needs to be monitored as closely as possible during production.

1.1.3.3.2. Complete stave failure modes should be clearly identified. The design should be analyzed in terms of these failure modes, trying to minimize their effect on overall performance. For instance, long daisy chains should allow break points, as the CDF Priority Bypass Chip solution, clock and control lines should have reasonable connection redundancy and what-if scenarios should be developed for foreseeable problems. The possible failure of the SVX4 chip in stressful conditions (such as high radiation rate, high temperature, etc.) should be examined and mitigated as much as possible.

1.1.3.3.3. The designs should be reviewed to eliminate single-point failures and to identify high-risk items such as connectors and couplings in the cooling system. These require special attention in long-term testing.

1.1.4. Quality assurance

1.1.4.1. Findings:

1.1.4.1.1. Both experiments base their approach to testing and burn-in on their experience with the Run IIa detectors.

1.1.4.1.2. Both experiments have developed and presented a preliminary production QC/QA plan.

1.1.4.2. Comments:

1.1.4.2.1. The planned burn-in at low temperature can be contrasted with the approach taken in space sciences to burn-in the hybrids at elevated temperatures to eliminate infant mortality, and to subject the finished but unpowered ladders or staves to conditions below the operational temperature to find solder-joint and glue-joint problems due to thermal stress.

1.1.4.3. Recommendations:

1.1.4.3.1. The groups should develop a comprehensive QC/QA document describing the tests to be done on each component and on the assemblies. They should take into account the industry standard procedures in terms of elevated temperature reliability testing. Particular care should be devoted to performing as extreme a test as possible on each component prior to assembly. These may include low and high temperature cycles, mechanical stress tests, elevated temperature burn-in, and should be focused on provoking failure early in the assembly process, thus reducing the rework rate and increasing the reliability.

1.1.5. Mechanical structure and cooling

1.1.5.1. Findings:

1.1.5.1.1. The two groups are pursuing different solutions for the stave cooling tube material: carbon fiber for D0 and PEEK for CDF. The difference is mainly motivated by the different aspect ratio of the channels, deriving from the different structure of the staves. Both solutions seem well justified and viable, although some technical issues remain. CDF has previous positive experience with PEEK.

1.1.5.1.2. D0's carbon fiber solution also improves the rigidity of the stave. A large amount of investigation on the characteristics of the carbon fiber has been done.

1.1.5.1.3. CDF made good progress fabricating the fixtures for the module assembly. D0 is ready to fabricate the fixtures, but doesn't yet have them in hand.

1.1.5.1.4. The collaborations have been given new guidelines as to the bunch structure and the luminosity goals for Run IIb. The new baseline with 396 ns bunch spacing yields a longer luminous region than the original 132 ns option (28 cm vs. 15 cm). This will lead to a reduced tracking efficiency for D0 due to the shorter inner layer staves.

1.1.5.2. Comments:

1.1.5.2.1. Each cooling pipe solution has its own merits and problems. The carbon fiber long-term stability is not well known. In addition, being conductive, it will induce noise on the sensors, requiring an effective grounding scheme to minimize the noise.

1.1.5.2.2. PEEK is intrinsically much less rigid than carbon fiber, although its stability and radiation resistance are well known. Square to round transition points and gluing may pose reliability problems with this material. A carbon fiber solution could reduce the gravitational sag of the stave and increase the vibration frequency of its fundamental mode. Although CDF claims that 200 μm sag does not affect SVT capability, it is still a concern to us.

1.1.5.2.3. CDF experienced failure of the cooling interconnects in the ISL. This points to the need for both experiments to engineer carefully their cooling system and to leave ample time for long-term testing.

1.1.5.3. Recommendations:

1.1.5.3.1. The long-term stability of the carbon fiber tubes should be confirmed through accelerated aging tests. For the PEEK tubes the reliability of the square to round transition and of the glue connection should be measured. Long-term leak tests under pressure and thermal stress should be carried out.

1.1.5.3.2. CDF should evaluate the relative advantages and disadvantages of the carbon fiber solution.

1.1.5.3.3. The committee recommends that high priority be given to the fabrication of a mechanical stave to characterize mechanical, cooling capability and robustness against thermal stress, as well as other mechanical properties.

1.1.5.3.4. D0 should conduct further studies on the impact of the 396 ns option and consider increasing the length of the inner layers to recover the efficiency loss.

1.1.6. Management and schedule

1.1.6.1. Findings:

1.1.6.1.1. The two collaborations attempted to reconcile and/or understand differences in their cost and manpower estimates. Nonetheless, D0's technical manpower estimate (130,000 hours) is about 60,000 hours higher than CDF's (70,000 hours).

1.1.6.1.2. D0 has shown a study of schedule sensitivity to slippage of intermediate tasks. Most task slippage has apparently little impact on the project end date.

1.1.6.2. Comments:

1.1.6.2.1. Both teams show adequate management structure, although it is not clear if the configuration control is adequate. D0 has all level 3 managers named and in place, while CDF has only subsystem managers.

1.1.6.2.2. The committee did not have enough time to examine the manpower estimate in details and to understand whether the differences and the absolute values are well justified.

1.1.6.2.3. The committee understands that the manpower available at Sidet during the silicon upgrades might not be sufficient to satisfy the needs of all the FNAL silicon projects.

1.1.6.2.4. The schedule sensitivity shown by D0 has only limited interest, since no leveling was applied, and likely the resources in the varied schedules are over-allocated.

1.1.6.3. Recommendations:

1.1.6.3.1. The two collaborations should continue to work on the budget and manpower comparison to identify the causes for the differences, and to justify them.

1.1.6.3.2. CDF should name the level 3 managers as soon as possible.

1.1.6.3.3. Schedule sensitivity should be analyzed by both experiments including manual or automatic leveling of resources.

1.1.6.3.4. The groups should define a clear process through which the design of parts is approved before starting production or procurement. This process may include sign-off procedures, final design reviews, production readiness reviews. This is urgent for those parts that need to be ordered soon. In some cases such procedures may already be in place.

1.1.6.3.5. The cooperation between the experiments on crucial technical issues should continue and be reinforced. Commonality of the two designs has already been well exploited but continuing collaboration during construction is crucial to ensure a timely completion of the project.

1.2. CDF Calorimeter Upgrades

1.2.1. Central Preradiator Upgrade

1.2.1.1. Findings:

1.2.1.1.1. The TDR describes convincingly the need for replacing the Central Preshower Detector. It also describes the structure of the proposed detector in about 1.5 pages. The TDR is accompanied by a 12-page WBS dictionary and twelve milestones are indicated in the Gantt chart for the task. The cost of this upgrade is estimated to be $700K plus 30% contingency, although costs given in the presentation were not the same as in the WBS dictionary which was made available on the web.

1.2.1.2. Comments:

1.2.1.2.1. The physics case for the upgrade appears to be well justified and the technical risks are minimal.

1.2.1.2.2. There is no indication of any technical drawings associated with the design. Because of this it is difficult to know how the phototubes are mounted and how the fibers are routed. No prototype studies of the baseline design are presented.

1.2.1.2.3. The TDR does not address the issue of the performance of the multi-anode PMTs in the return field of the solenoid magnet. In discussion, however, the proponents indicate that they are already operating these PMTs in similar field conditions.

1.2.1.2.4. Many tasks in the WBS dictionary have no time duration. For example, the R&D task requires $101K but is completed in one day. There is no indication of the deliverables from the R&D work.

1.2.1.3. Recommendations:

1.2.1.3.1. This task would be strengthened by more evidence of engineering work in planning the design. The TDR indicates that prototype studies using the Dubna scintillator are underway. Any quantitative preliminary results from this work would be useful.

1.2.1.3.2. A statement should be added to the TDR on experience with these phototubes in similar magnetic field conditions.

1.2.1.3.3. The costs given in the WBS dictionary should be reconciled with those given in the presentation.

1.2.1.3.4. This task is ready to baseline, although the documentation would benefit from the additions just mentioned.

1.2.2. EM Calorimeter Timing

1.2.2.1. Findings:

1.2.2.1.1. The TDR describes the need for a timing measurement from the 960 PMTs of the central EM calorimeter and 768 PMTs of the plug EM calorimeter. The technical aspects of the work are described in about 1.5 pages, There is a clear description of the inductive splitter, its negligible effect on the calorimeter energy measurement, and the time resolution obtained using an LED signal. The TDR is accompanied by a 13 page WBS dictionary. The total cost of this upgrade is estimated to be $250K plus 30% contingency.

1.2.2.2. Comments:

1.2.2.2.1. The scope of this task appears to be clearly defined and a working solution has been demonstrated.

1.2.2.2.2. This does not appear to be an upgrade of the highest priority but it would certainly strengthen the characteristics of the detector.

1.2.2.3. Recommendations:

1.2.2.3.1. This task appears ready to baseline.

1.3. CDF Trigger/DAQ Upgrades

1.3.1. General Comments on CDF Trigger Upgrades

1.3.1.1. Findings:

1.3.1.1.1. The Run IIa CDF trigger design is sound and presents a well defined upgrade path for Run IIb operation.

1.3.1.1.2. The maximum rates at each of the three trigger levels in the proposed Run IIb system are as follows:

- Level-1 accept: ~50kHz
- Level-2 accept: ~1kHz
- Level-3 accept: ~85Hz

1.3.1.1.3. Extrapolations performed by the CDF trigger group indicate that the proposed Run IIb trigger system will function well at a luminosity of 4×10^{32} /cm^2/sec with 396 ns bunch spacing.

1.3.1.2. Comments:

1.3.1.2.1. The CDF detached vertex trigger has demonstrated impressive performance under the present Run IIa conditions.

1.3.1.2.2. The CDF scheme for incremental installation and testing is commendable.

1.3.1.2.3. If a luminosity of 2×10^{32} /cm^2/sec is reached at the end of Run IIa, the present CDF trigger and DAQ systems may be operating near or beyond their design capacity.

1.3.1.2.4. Commissioning the proposed Run IIb trigger system will require significant manpower resources as well as careful co-ordination between the trigger group and all other detector components, most notably data acquisition. Although a general installation plan is in place, careful consideration should be given to the details of the commissioning effort as the Run IIb projects evolve.

1.3.1.3. Recommendations:

1.3.1.3.1. Estimates of various trigger line rates for Run IIb conditions use current performance figures extrapolated linearly to higher luminosity. To verify these extrapolations, the CDF trigger group should also consider using simulated events for high luminosity conditions which include multiple interactions per bunch.

1.3.2. XFT Upgrade

1.3.2.1. Findings:

1.3.2.1.1. The CDF Run IIb Level-1 track trigger, which the proponents call the eXtremely Fast Tracker (XFT), is an upgraded version of the track trigger designed for Run IIa. The new design preserves the general architecture of the Run IIa track trigger.

1.3.2.1.2. Since roughly 50% of all physics triggers involve tracks from the XFT, it is crucial to overall system performance.

1.3.2.1.3. The upgraded XFT derives its main performance gain by exploiting a factor-of-three improvement in the precision of the timing information transmitted from the upgraded TDCs to the XFT. This reduces the effective size of the trigger elements, leading to a reduction in fake tracks and an improvement in the P_t and φ_0 resolution.

1.3.2.1.4. The system also incorporates stereo information, which further reduces the rate for fake lepton triggers. This feature would be particularly crucial if the collider were to run at 132 ns, which would eliminate the factor-of-three timing advantage cited above.

1.3.2.1.5. The overall effort required for the XFT is significant, involving roughly 70 complex PC boards and at a total estimated cost of \$2.6M dollars. Significant effort has already been expended on many design details of the XFT. However, it appears that a considerable amount of work remains on printed circuit board design and firmware development.

1.3.2.2. Comments:

1.3.2.2.1. A successful upgrade of the XFT track trigger is essential to successful operation of CDF during Run IIb, as has been presented by the proponents. Their simulations show a rapid deterioration in P_t and φ_0 resolution as well as a rapid increase in fake rate as the number of interactions per crossing increases.

1.3.2.2.2. The XFT design appears to be sound and is sufficiently well advanced to be credible. Although the committee did not examine costs and technical risks in detail, the proposed design appears to be within the state of the art and is sufficiently well defined to be reliably costed. There is some indication from comparison with similar systems that the cost estimates may be on the high side.

1.3.2.2.3. Complete confidence in the design can only be gained through use with realistic signals from the detector.

1.3.2.3. Recommendations:

1.3.2.3.1. The proponents should further validate the upgraded XFT design by studying its performance using a software emulation based on Run IIa data and/or (if possible) by testing prototype boards with actual signals from the detector.

1.3.2.3.2. The committee feels this task is ready to baseline.

1.3.3. TDC Replacement for the Central Outer Tracker

1.3.3.1. Findings:

1.3.3.1.1. The current TDC modules used for the COT will seriously limit the ability of the CDF detector to take data at Run IIb rates for several reasons:

- Hit processing is performed only after a Level-2 accept, hence the total processing plus readout deadtime associated with a Level-2 accept is too large.
- The readout of the TDC buffers via VME block transfer is too slow.
- Data transfer out of the TDC crates via TAXI is too slow.

1.3.3.1.2. The current modules do not provide the information needed by the proposed XFT upgrade.

1.3.3.1.3. The proposed TDC upgrade provides an elegant solution to all of these problems. The time to digital conversion, hit processing, buffering, and readout can be implemented in a single Altera Stratix FPGA.

1.3.3.1.4. The cost of this task is approximately $1.67M, which includes 33% contingency.

1.3.3.2. Comments:

1.3.3.2.1. The committee feels the CDF trigger group is pursuing the correct solution to this problem and that the proposed upgrade should be implemented.

1.3.3.3. Recommendations:

1.3.3.3.1. The committee feels this project is ready for baselining.

1.3.4. Level-2 Decision Crate

1.3.4.1. Findings:

1.3.4.1.1. The estimated cost for this item is $215K plus 30% contingency for M&S.

1.3.4.1.2. The proposed replacement of the Level-2 decision crate addresses the need for increased bandwidth and computing power for the Level-2 trigger system. It is proposed to replace the Level-2 systems by newly developed electronic boards, which have been developed in the context of the Level-2 test-stand (Pulsar system).

1.3.4.1.3. The proponents list additional reasons to do this upgrade, pointing out that the currently employed alpha processors do not provide a viable hardware platform for the longer term, and that the current diversity of interface boards to the front end electronics poses serious maintenance issues.

1.3.4.1.4. The proposed new boards have a common design for all systems, with small interface boards specific to each detector component. The design makes use of standard commodity PCs to provide CPU power, instead of using embedded CPUs. It also uses the S-LINK bus system developed and implemented at CERN and used by Atlas.

1.3.4.2. Comments:

1.3.4.2.1. This approach appears reasonable. It provides a simplification of the system and a clear upgrade path for the increased needs for processing power.

1.3.4.2.2. However, no specific study was presented justifying the specific design and performance. The documentation provided in the TDR is very abbreviated and involves less than a page of text. It is beyond the scope of this review to look in detail into the design, implementation, costs and risks.

1.3.4.3. Recommendations:

1.3.4.3.1. Because of the limited information available a review of the technical solution, cost estimate, and schedule still remains to be done. The committee recommends that the project provide a more detailed report in the future. For the time being, the resources for this upgrade should become part of the "project trust fund" recommended by the PAC.

1.3.5. Event Builder Switch

1.3.5.1. Findings:

1.3.5.1.1. The cost estimate for this item is \$414K + 30% contingency. The estimate only considers the equipment costs to upgrade existing hardware and to provide spares.

1.3.5.1.2. The upgrade is to increase the bandwidth of the system and hence to accommodate the higher rates and larger event sizes of Run IIb operation. The required bandwidth is estimated to be at least 250 Mbytes/second.

1.3.5.1.3. The current CDF Event-Builder consists of a 32 port OC3 ATM switch with a bandwidth limit of 240 Mbytes/s. About 60% of that bandwidth has been achieved in benchmark tests using simulated event sizes of Run IIb. It is possible that up to 80% of the theoretical limit could be achieved after tuning the system.

1.3.5.1.4. CDF proposes to upgrade the switch to provide OC12 ATM links. This would increase the throughput of the system to a theoretical limit to 1 Gbyte/s.

1.3.5.2. Comments:

1.3.5.2.1. The committee did not hear any evaluation of alternatives to the proposed solution, such as replacing the ATM switch with a Gigabit Ethernet switch. D0 has successfully implemented such a system, and this approach may be feasible, cost effective and would remedy issues with the ATM interfaces on the processing nodes.

1.3.5.2.2. It is noted that the OC12 (and OC3) interfaces on Linux computers are not commodity items. The development and integration of these drivers and their upgrade to accommodate new versions of Linux requires a high level of expertise. CDF has currently decided to freeze the Linux version on their event builder nodes because of incompatibilities of the OC3 interface drivers with the latest kernel versions. It will be important to keep open the option of upgrading to more recent versions of the kernel if one wishes to be able to use modern higher performance hardware.

1.3.5.2.3. The committee feels that this poses a significant risk to the project which has not yet been adequately addressed. There is a possible need for a significant software effort to modify and integrate drivers for OC12 which has not been accounted for in the project costs. The group proposes using students for this work. Since individual students only remain with the project for a limited time, it will be important to ensure that the code they produce can be maintained by others.

1.3.5.3. Recommendations:

1.3.5.3.1. The committee recommends an explicit assessment of this issue. The risks related to these issues should be elaborated and a mitigation plan proposed.

1.3.5.3.2. All related efforts and costs, including software and integration, should be tracked by the project, and thus be included in the WBS and schedule (WBS item 1.3.4.1), even if some of the effort is entered as zero-cost items.

1.3.5.3.3. This task is ready to baseline but before procurement other technical solutions should be examined.

1.3.6. Level-3 Processor Farm

Findings, comments and recommendations in these paragraphs concern both D0 and CDF, and are presented here for both experiments.

1.3.6.1. Findings:

1.3.6.1.1. The Run II Level-3 systems of D0 and CDF are scalable farms of Linux PCs, allowing the experiments to make use of commodity hardware for compute nodes, networking infrastructure and data storage. The committee commends D0 for their very successful effort in bringing their commodity hardware Level-3 system into operation.

The designs for the Level-3 systems allow a straight-forward upgrade to increase the throughput and processing power. Replacing older compute nodes with new higher-performance commodity hardware will take advantage of Moore's law to obtain the required performance increase. The committee in general agrees with the need for regular upgrading the systems during Run IIa. This will provide the necessary computing power at the start of Run IIb.

The estimated computing needs for Run IIb are based on a linear scaling of the current processing needs to the Run IIb situation with multiple interactions per bunch crossing. CDF and D0 are starting from quite different processing needs. This leads to the estimated Run IIb requirement of 6 CPU seconds per event for CDF, and 1.5 CPU seconds per event for D0, on a 1 GHz Pentium III. There is the assumption that CPU performance will continue to increase by a factor of about 1.7 each year.

CDF proposes to arrive at the required level of performance by upgrading 85 nodes each year in FY03, FY04 and FY05, with estimated costs of $390K plus 30% contingency.

D0 proposes to upgrade 32 nodes in FY04 and 64 nodes in FY05, at a total cost of $210K plus 70% contingency.

1.3.6.1.2. CDF does not foresee any upgrade of other DAQ-related computing systems as part of the scope of this project, although they will certainly be needed. The committee was informed that CDF considers those costs to be part of regular computing upgrades funded as operating expenses for Run IIa.

D0, on the other hand, proposes an upgrade of DAQ-related computing systems, including data base servers and data storage servers. The cost of these for D0 is $247K plus 50% contingency.

1.3.6.2. Comments:

1.3.6.2.1. The committee notes that if both experiments targeted their Level-3 upgrades solely for Run IIb, they would probably procure all processors as late as possible. That would allow CDF to obtain 1.8 times the performance at equal costs, or to decrease the costs by 60% with the corresponding benefit of being able to reduce the size of the system.

1.3.6.2.2. This committee was unable to look in detail at the proposed technical solutions, validity of approach and estimated costs. It feels that D0 estimated costs for the “host systems” upgrade are relatively high for providing a rather moderate, although highly available storage system of about 5TB and two data base servers. Those costs, like the costs for the farm upgrades, could probably be lower if the upgrades were targeted to 2006, instead of being available already in FY04.

1.3.6.2.3. The committee does not disagree with the claim that computing upgrades in the DAQ area will be needed already for the expected increase in luminosity and rates of Run IIa.

1.3.6.2.4. The use of commodity systems for Level-3 and DAQ has resulted in large similarities between the computing systems used for online and offline. The expected rise in Level-3 output rates and the increase in event size because of higher detector occupancies will have an important impact on the need for offline computing and data handling systems. These needs include network throughput and physics analysis resources at outside institutes. The increased resource requirements will not be a step function with the start of Run IIb but will rise progressively during Run IIa as luminosity increases.

1.3.6.3. Recommendations:

1.3.6.3.1. The committee feels that the experiments have not yet created a plan reconciling both the need for upgrades during Run IIa and the provision of computing power for Run IIb. It recommends developing such an integrated plan for computing upgrades taking into account both needs and thereby optimizing the use of resources.

1.3.6.3.2. The committee recommends that software and computing issues both in online (DAQ and Level-3) and in offline be addressed by a separate standing Run IIa/b computing review committee.

1.4. CDF Installation

1.4.1. Findings:

1.4.1.1. The present installation plan is described in about 2/3 of a page in the TDR. An installation time of 34 weeks is called for, including 50% contingency. The task is supported by a 27 page WBS dictionary. The estimated cost is $768K plus a contingency of $502K. The work requires an average of 17 FTEs over its duration and is based on a 40-hour week.

1.4.2. Comments:

1.4.2.1. The installation manager and his team are highly experienced and appear to be fully able to organize and complete the work. The present description in the TDR, however, is quite abbreviated. The presentation to the committee was much more informative and complete.

1.4.2.2. No Level 3 managers are identified for this task.

1.4.2.3. No profile of the manpower is shown over the duration of the task.

1.4.2.4. Planning is needed of the ramp-up of the installation process. An effort should be made to minimize simultaneous responsibility of individuals for both finishing the construction and planning installation.

1.4.3. Recommendations:

1.4.3.1. It would strengthen the TDR to include a summary of the tasks to be done and a monthly breakdown of the manpower required, according to type.

1.5. D0 Level-1 Trigger Upgrades

1.5.1. General Comments: related to the D0 Trigger

1.5.1.1. Findings:

1.5.1.1.1. The maximum rates at each of the three trigger levels in the proposed Run IIb system are as follows:

- Level-1 accept: ~5kHz
- Level-2 accept: ~1kHz
- Level-3 accept: ~50Hz

1.5.1.1.2. For D0 the acceptable rate from Level-1 is approximately an order of magnitude lower than for CDF.

1.5.1.2. Comments:

1.5.1.2.1. The Run IIa D0 trigger design appears able to meet its technical specifications and represents a reasonable basis for the Run IIb upgrade.

1.5.1.2.2. Despite this, if a luminosity of 2×10^{32} /cm^2/sec with 396 ns bunch spacing is reached near the end of Run IIa, the present D0 trigger and DAQ systems may be near the limits of their design capacity.

1.5.1.2.3. Simulations and extrapolations performed by the D0 group indicate that the proposed trigger upgrade for Run IIb will function well at a luminosity of 2 x 10^{32} and 396 ns. The committee accepts their conclusion.

1.5.1.2.4. Early deployment of some of the upgrades could help late in Run IIa if the luminosity is high.

1.5.1.2.5. The D0 scheme for incremental installation and testing is commendable.

1.5.1.2.6. Studies for a luminosity of 4×10^{32} /cm^2/sec and a bunch spacing of 396 ns indicate that some of their trigger components, most notably the high P_t track trigger, may again be very close to their operational limit. This represents potential scope risk and needs to be examined further. In particular, the performance of the Level-2 silicon tracker, which is an important part of their trigger, needs to be simulated under these conditions. Since the issue is associated with headroom beyond the baseline luminosity we believe it should not impact baselining the project. The contingency might reflect this risk.

1.5.1.3. Recommendations:

1.5.1.3.1. The D0 Level-2 and Level-3 trigger algorithms for Run IIb have not yet been finalized. Some of this software will be developed during Run IIa, but substantial additional effort will be needed for Run IIb. We recommend that an explicit plan be developed for producing the Level-2 and Level-3 trigger algorithms and associated software tools.

1.5.2. Level-1 Tracking Trigger

1.5.2.1. Findings:

1.5.2.1.1. The Run IIb Level-1 track trigger, which employs hits from the charged fiber tracker (CFT), is an upgraded version of the L1CTT trigger designed for Run IIa.

1.5.2.1.2. Since almost all physics triggers involve the combination of information from another detector subsystem with tracks from the L1CTT, it is crucial to overall system performance.

1.5.2.1.3. The new design preserves the general architecture of the Run IIa design and derives its main performance gain from the use of single fiber hits, as opposed to the doublets employed in the Run IIa design. Implementation of this logic involves a substantial amount of new hardware (M&S costs of $1.1M, including contingency) primarily for replacing the 80 daughter boards that implement the L1CTT logic.

1.5.2.1.4. Significant effort has already been expended on many design details of the upgraded L1CTT. For example, the group is well advanced in defining the logic and establishing the capacity of the field programmable gate arrays required to implement it. Results of reasonably detailed physics simulations were presented which showed improvements in rejection of more than an order of magnitude relative to Run IIa.

1.5.2.2. Comments:

1.5.2.2.1. A successful upgrade of the L1CTT track trigger is essential to successful operation of D0 during Run IIb. This comment is supported by simulation data presented by the proponents, which show a rapid degradation in performance of the current system even at modestly increased occupancies, such as those that may be encountered near the end of Run IIa.

1.5.2.2.2. The L1CTT design appears to be sound. The design is sufficiently well advanced to be credible. Additional confidence would come from more detailed studies based on the superposition of real minimum bias events.

1.5.2.2.3. Although the committee did not examine costs and technical risks in detail, the proposed design appears to be within the state of the art and is sufficiently well defined to be reliably costed.

1.5.2.3. Recommendations:

1.5.2.3.1. The proponents should further validate the upgraded L1CTT design by studying its performance using a software emulation based as closely as possible on Run IIa data but with multiple events superimposed as expected for the conditions of Run IIb.

1.5.2.3.2. This subsystem appears ready to baseline.

1.5.3. Level-1 Calorimeter Trigger

1.5.3.1. Findings:

1.5.3.1.1. The D0 Level-1 calorimeter trigger is based on signals from 1280 EM towers and 1280 hadronic towers, each 0.2×0.2 in η and φ.

1.5.3.1.2. The input signals to the trigger are rather slow, with a 150 ns rise time and 400 ns width. This makes their association with a given bunch crossing difficult for a 132 ns bunch crossing time. Operation at 396 ns is more straightforward.

1.5.3.1.3. The calorimeter tower size for the trigger is currently much smaller than the characteristic size of a hadronic jet. This leads to a very slow turn on of the jet trigger efficiency as a function of E_t. For example, to obtain 100% efficiency for a 60 GeV jet requires a tower threshold of only 6 GeV. The result is a Level-1 jet trigger dominated by low energy jets. The efficiency for electrons and photons is similarly degraded for impact points near tower boundaries. The proposal is to implement a sliding window algorithm and a larger jet tower size to sharpen the trigger threshold.

1.5.3.1.4. The center of a jet is estimated by a sliding window of 0.4×0.4 over the trigger towers to locate local maxima. The jet energy is estimated by summing over a region of 0.8×0.8 centered on a maximum found by the sliding window.

1.5.3.1.5. The result is that for a setting which gives 85% efficiency for jets above 40 GeV, the trigger rate is reduced by a factor of 3. The electron trigger is similarly strengthened. The performance for particular channels of Higgs production is shown.

1.5.3.1.6. The design includes the capability to add signals from the inter-cryostat detectors into the energy trigger to further improve resolution. It also introduces the possibility of enriching the trigger in τ leptons through the presence of a very narrow jet in the calorimeter.

1.5.3.1.7. The TDR contains 83 pages describing the principles of operation, the performance, and the implementation details of this upgrade. The schedule is described in a 92-line Gantt chart with 9 high-level milestones. The cost of this upgrade is $1.3M, including 43% contingency.

1.5.3.1.8. Saclay, Columbia, and Michigan State propose to take the lead responsibility for the task.

1.5.3.2. Comments:

1.5.3.2.1. The TDR contains extensive detail on studies done to explore and optimize the performance of the system, as well as on the design. The proponents would be well served if they could characterize, relative to the Run IIa trigger, the improvement brought by this upgrade to the overall significance of the Higgs signal. Other global performance figures would also be helpful to make clear the impact of this upgrade.

1.5.3.2.2. The proponents have indicated in response to questioning that elimination of the digital filter, which is less critical for 396 ns operation, would save less than $50K since FPGAs are needed in any case to format the data for the TAB boards which perform the sliding window calculation.

1.5.3.2.3. The three principal institutions all have extensive experience with complex trigger systems.

1.5.3.2.4. The proponents indicate that they plan to test parts of the new system during Run IIa using signals from the present detector. We view this as a very valuable process.

1.5.3.3. Recommendations:

1.5.3.3.1. The proponents should try to characterize the performance of the upgraded system with a few global figures of merit. The PAC has emphasized the Higgs detection significance.

1.5.3.3.2. This task appears ready to baseline.

1.5.4. Calorimeter-Track Matching Trigger

1.5.4.1. Findings:

1.5.4.1.1. The high rate of fake tracks and showers becomes a problem as luminosity increases. D0 studies have shown that at high luminosity the proposed Level-1 track-shower matching system will reduce the rate of false medium P_t electron triggers by a factor of two or three. It can also be used to reject fake tracks by a factor of one to two orders of magnitude.

1.5.4.1.2. The proposed Cal-Track design combines information from the upgraded CFT track trigger as well as the upgraded calorimeter trigger to correlate in φ hits between the two. The proposed design uses the fact that an eight-fold increase in φ granularity will be available from the proposed calorimeter upgrade.

1.5.4.1.3. The system exploits the existing design for a similar system used to correlate CFT tracks with hits in the muon system. The use of an existing design minimizes both the cost and risk of the proposed upgrade.

1.5.4.1.4. The cost of this project is approximately $260K, which includes 31% contingency.

1.5.4.2. Comments:

1.5.4.2.1. While the D0 collaboration has not explicitly made the case that the Cal-Track project is needed in order to successfully pursue a high P_t physics program, the committee feels that this upgrade is a prudent and cost-effective measure given that the overall trigger system may be struggling to provide adequate rejection at the highest Run IIb luminosities. In particular, this system could be a key ingredient in keeping the rate of high P_t track triggers to a tolerable level at luminosities above 2×10^{32} /cm^2/sec, where 6 or more minimum bias events are expected from each crossing and fake track trigger rates are a potential problem.

1.5.4.3. Recommendations:

1.5.4.3.1. The committee feels this project is ready to baseline.

1.6. D0 Level-2 Trigger Upgrades

1.6.1. Level-2 Beta Trigger

1.6.1.1. Findings:

1.6.1.1.1. The presented project cost is $64K including 30% contingency. No Fermilab labor has been assigned to this item.

1.6.1.1.2. The system will already be commissioned for Run IIa, where it will replace the current Level-2 alpha boards. The group is expecting to obtain pre-production Level-2 beta boards with current-generation (commercially available) processor boards for Run IIa. They will commission the system with about 26 boards this year, completely replacing the existing Level-2 alpha boards.

1.6.1.1.3. The specified costs for this project are solely for upgrading 12 of the CPU boards to provide increased processing power to the Level-2 trigger.

1.6.1.2. Comments:

1.6.1.2.1. Although the Level-2 computing boards were a high-risk item for Run IIa, the committee feels that this project presents only moderate risk for Run IIb.

1.6.1.2.2. It is, however, of central importance for the D0 upgrade to achieve the goals for the Level-1 and Level-2 output rates. This will require the rejection of substantially increased backgrounds, specifically from the tracking triggers. With the Level-1 upgrade many of the current Level-2 cuts will be moved to the Level-1 trigger. D0 will need to develop a new set of Level-2 algorithms to keep the Level-2 output rate below 1 kHz.

1.6.1.2.3. General ideas for revised Level-2 algorithms were presented, such as moving the vertex finding to Level-2. The exact effectiveness of these cuts over Level-1 will need to be studied. The required increase in CPU performance needs to be estimated and a method devised to parallelize the processing on several nodes.

1.6.1.2.4. The committee finds that most of the effort in this item is in providing the necessary physics algorithms on the Level-2 processors. This effort is not spelled out in the project, but its success is essential for the success of the Level-2 project. Currently the Level-2 software effort consists of a reasonable sized group of 6-8 physicists. The project will need to track this effort.

1.6.1.3. Recommendations

1.6.1.3.1. Upgrading the single board computers as this task proposes represents a clear path towards obtaining the required processing power. However, the committee has not seen detailed studies on what resources are needed to obtain the required cut in Level-2 rate. Thus the cost estimates, which foresee replacing 12 of the boards for Run IIb, should be considered somewhat preliminary, and contingency should be foreseen in case more CPU resources are needed.

1.6.1.3.2. The committee would like to see a Level-2 trigger report as a milestone, where simulation studies and tests are compiled to show the rejection power of the Level-2 for Run IIb running. The report should address the required processing power and bandwidth, and outline a plan for providing the required Level-2 software.

1.6.1.3.3. The committee feels that the proposed solution is reasonably straight forward and cost effective and is ready to baseline.

1.6.2. Level-2 Silicon Track Trigger

1.6.2.1. Findings:

1.6.2.1.1. The Silicon track trigger (STT) is a Level-2 trigger preprocessor that combines information from the silicon microstrip tracker and the Level-1 fiber tracker and produces high resolution momentum and impact parameter information for each track candidate. Since information is correlated between two independent detector systems the rate of fake tracks is also reduced. The cost of this project is approximately $329K including 43% contingency.

1.6.2.2. Comments:

1.6.2.2.1. The STT is a key component of the D0 trigger system. It plays an important role in all physics trigger lines and is the key ingredient in lines using detached vertices to tag b-jets.

1.6.2.2.2. Since the Run IIa version of the STT will be commissioned during the next six months, verifying that the present device functions within design specifications should be a key milestone in the execution of the Run IIb system.

1.6.2.3. Recommendations:

1.6.2.3.1. The performance of the STT trigger at the highest proposed Run IIb luminosities with a bunch spacing of 396 ns should be studied further using the same detailed simulations used to validate the design for operation with a bunch spacing of 132 ns. These studies, combined with the experience from commissioning the STT for Run IIa, should be used to guide the Run IIb STT project.

1.6.2.3.2. The committee feels this project is ready to baseline.

1.7. D0 DAQ/Online Upgrades

See comments under 1.3.6 where these issues are discussed for both CDF and D0.

1.8. D0 Installation

1.8.1. Findings:

1.8.1.1. The main installation work involves replacing the D0 silicon. This operation involves steps similar to those carried out during the Run IIa installation and therefore can be planned with reasonable certainty. During this operation the detector will remain on the beamline, unlike the situation for silicon installation in Run IIa.

1.8.1.2. The installation team estimates that a total of 30 weeks will be required from the time the Tevatron stops until the detector is closed. The cost of the effort is $1.3M, including contingency, and a team of 45 (peak) physicists, engineers, and technicians. The average manpower requirement is 24 FTEs.

1.8.2. Comments:

1.8.2.1. An appropriately detailed and credible plan was presented for the installation. A management team has been named. Although the current installation team is somewhat understaffed, the proponents argue that additional experienced manpower will become available as the construction part of the project winds down.

1.8.3. Recommendations:

1.8.3.1. The group should revisit the installation plan as the date approaches and the construction efforts are completed. At that time the individuals available will be clearer. It will be important to minimize additional responsibilities of the management personnel trying to complete construction tasks.